

Abstract

A database system that can synchronize all or a part of its contents over a limited bandwidth link is described. The lowest layer of the system, the bedrock layer, implements a transactional block store. On top of this is a B+-tree that can efficiently compute a digest (hash) of the records within any range of key values in $O(\log n)$ time. The top level is a communication protocol that directs the synchronization process such that minimization of bits communicated, rounds of communication, and local computation are simultaneously addressed.